

APPLICATION FOR UNITED STATES LETTERS PATENT

FOR

**A Scalable Directory, Master and Slave Server Based
Architecture For Hosting Application Services**

Inventor(s):

**Matthew F. Hillman
Christopher O. Peterson
John Y. Geer**

Prepared by:

**COLUMBIA IP LAW GROUP, LLC
WASHINGTON OFFICE**

"Express Mail" label number EL743034535US

Deposit Date: March 16, 2001

0934426-034604
FOOTED 922F860

A Scalable Directory, Master and Slave Server Based Architecture
For Hosting Application Services

Related Application

[0001] This application is a non-provisional application of the earlier filed provisional application number 60/252,677, filed on November 20, 2000, and claims priority to the earlier filed '677 provisional application, whose specification is hereby fully incorporated by reference.

BACKGROUND OF THE INVENTION

1. Field of the Invention

[0002] The present invention relates to the field of hosting application services. More specifically, the present invention relates to the issue of hosting application services in a scalable manner to allow the hosting resources to be gradually, but easily added, as the services, instantiations and/or user bases of the services increase. The present invention has particular application to hosted multi-party packet based talk radio shows (commonly known as hosted Internet Talk Radio).

2. Background Information

[0003] With advances in integrated circuit, microprocessor, networking and communication technologies, increasing number of devices, in particular, digital computing devices, are being networked together. Devices are often first coupled to a local area network, such as an Ethernet based office/home network. In turn the local area networks are interconnected together through wide area networks, such

as ATM networks, Frame Relays, and the like. Of particular notoriety is the TCP/IP based global inter-networks, the Internet.

[0004] As a result this trend of increased connectivity, increasing number of applications that are network dependent are being deployed. Examples of these network dependent applications include but are not limited to, email, net-based telephony, world wide web, various types of e-commerce and Internet Talk Radio Shows. An "Internet Talk Radio Show" is an online form of communication or entertainment in which a host, using his computer, speakers, and a microphone, broadcasts an audio show to a variety of listeners over the Internet, who listen to the show through their computer and speakers, and can interact with the host through a microphone. The intended experience for the participants is to be similar to conventional terrestrial talk radio, however the technology involved bears no relationship to the technology of terrestrial radio. Further, increasing number of software applications that were traditionally licensed or distributed through discrete distribution medium, such as diskettes, CDRoms and the like, are being distributed online or offered as web based applications, through private intranets or public networks like the Internet.

[0005] The increased popularity of network based applications and offering of traditional software applications as application services have in turn led to the emergence of application service providers who host application services for their developers, to relieve the developers from having to maintain and operate the underlying infrastructure. A critical issue to the success of hosting application services is the ability of the application service providers to limit the deployment of the resources required to host the application services, and to gradually grow these

resources as the services, the instantiations and/or user bases of the services expand. However, the architecture must allow the newly required resources to be easily integrated with the deployed resources, preferably in a manner that is transparent to the users of the hosted application services. Accordingly, a highly flexible and scalable architecture to host application services is desired.

FOR E.O. 13526

SUMMARY OF THE INVENTION

[0006] A number of directory servers non-exclusively engage a first plurality of servers to serve as master servers for hosting one or more application services, and a second plurality of servers to serve as slave servers for facilitating clients accesses to the one or more hosted application services. During operation, the directory servers receive requests from clients to access a selected one or selected ones of the application services. In response, the directory servers select and assign selected ones of their non-exclusively engaged master servers to host the one or more accessed application services as well as select and assign selected ones of their non-exclusively engaged slave servers to facilitate the clients to access the one or more hosted application services.

[0007] In one embodiment, the selections and assignments of master and slave servers are performed in accordance with a number of metrics that reflect the ability of the master and slave servers to further undertake hosting of additional instantiation of the application services or additional application services, and facilitating accesses to the hosted application services, in view of the hosting and facilitating already undertaken. In one embodiment, the directory server periodically polls their non-exclusively engaged master and slave servers for operational information related to these decision metrics.

[0008] In one embodiment, the hosted application service is a hosted Internet Talk Radio System and its hosted multi-party talk radio shows. The clients may be hosts, callers or listeners of the talk radio shows. The master servers are

responsible for mixing the audio streams for the listeners, and routing the mixed and non-mixed audio streams to the listeners and callers accordingly.

FOR THE REASON

BRIEF DESCRIPTION OF DRAWINGS

The present invention will be described by way of exemplary embodiments, but not limitations, illustrated in the accompanying drawings in which like references denote similar elements, and in which:

Figure 1 illustrates an overview of the present invention, including the directory, master and slave servers, in accordance with one embodiment;

Figure 2 illustrates the method of the present invention, in accordance with one embodiment;

Figure 3 illustrates an overview of an application of the present invention to Internet Talk Radios;

Figure 4 illustrates the method of the present invention as applied to Internet Talk Radios; and

Figure 5 illustrates an internal component view of a computer system suitable for use to practice the present invention, in accordance with one embodiment.

DETAILED DESCRIPTION OF THE INVENTION

[0009] In the following description, various aspects of the present invention will be described. However, it will be apparent to those skilled in the art that the present invention may be practiced with only some or all aspects of the present invention. For purposes of explanation, specific numbers, materials and configurations are set forth in order to provide a thorough understanding of the present invention. However, it will also be apparent to one skilled in the art that the present invention may be practiced without the specific details. In other instances, well known features are omitted or simplified in order not to obscure the present invention.

[0010] Parts of the description will be presented in terms of operations performed by a processor based device, using terms such as data, tables, accepting, determining, inquiring, notifying, caching, routing, loading, and the like, consistent with the manner commonly employed by those skilled in the art to convey the substance of their work to others skilled in the art. As well understood by those skilled in the art, the quantities take the form of electrical, magnetic, or optical signals capable of being stored, transferred, combined, and otherwise manipulated through mechanical and electrical components of the processor based device; and the term processor include microprocessors, micro-controllers, digital signal processors, and the like, that are standalone, adjunct or embedded.

[0011] Various operations will be described as multiple discrete steps in turn, in a manner that is most helpful in understanding the present invention, however, the order of description should not be construed as to imply that these operations are necessarily order dependent. In particular, these operations need not be performed

in the order of presentation. Further, the description repeatedly uses the phrase “in one embodiment”, which ordinarily does not refer to the same embodiment, although it may.

Overview

[0012] Referring now to **Figure 1**, wherein a block diagram illustrating an overview of the present invention, including the directory, master and slave servers, in accordance with one embodiment, is shown. As illustrated, a number of directory servers **102**, a number of master servers **104**, and a number slave servers **106** are scalably coupled to each other in accordance with the present invention to host one or more application services and facilitate accesses to these one or more hosted application services by clients **108**. As will be readily apparent from the description to follow, servers **102**, **104** and **106** are advantageously coupled together in a scalable manner to allow only a necessary amount of these resources to be deployed to host the one or more application services, and additional resources be gradually and easily added as the application services, the instantiations and/or user bases of the hosted application services grow. For ease of understanding, the term “application service” will be used in its plural form, however unless otherwise specifically stated, each usage is to include its singular form. That is, the present invention may be practiced to host one or more application services.

[0013] Master servers **104** are employed to host the application services of interest, which may be any one of a variety of application services known in the art, including but are not limited to an Internet Talk Radio System, and its hosted multi-party talk radio shows. Master servers **104** host application services as instructed by directory servers **102**. Master servers **104** are non-exclusively engaged with the

directory servers **102** to host these application services, to be described more full below.

[0014] Slave servers **106** are employed to facilitate accesses by clients **108** to the application services hosted by master servers **104**. That is, all communications or interactions between the hosted application services on master servers **104** and clients **108** are made through slave servers **106**. Slave servers **106** are informed of the hosts of the application services (i.e. master servers **104**) as well as clients **108**, the slave servers **106** are to facilitate accesses. The information is provided to slave servers **106** by directory servers **102**. Similar to master servers **104**, slave servers **106** are not exclusively engaged with the directory servers **102** to facilitate accesses to the hosted application services, also to be described more full below.

[0015] Clients **108** request for accesses to the hosted application services with directory servers **102**, and are informed of the slave servers **102** through which clients **108** are to make the accesses. Depending of the type of application services, such as the earlier mentioned Internet Talk Radio System, clients **108** may be divided into classes and their accesses to the hosted application services may be facilitated in accordance with their class memberships. For example, for the Internet Talk Radio application, access to a talk radio show of by a client **102** may be facilitated depending on whether the client is a host of the talk radio show, a listener of the talk radio show or a caller of the talk radio show. [Note the term "host" is used in at least two contexts, a first context where a master server **104** "hosts" an application service, and a second context where a client **108** "hosts" a talk radio show (which itself programmatically speaking is "hosted" by a master server **104**).]

[0016] Except for logic to be provided to the various servers **102-106** and clients **108** to enable these devices to cooperate and interact with one another to enable the relatively easy and scalable expansion and contraction of server resources in proportion to the amount of application services and clients **108** supported, servers **102-106** and clients **108** are otherwise intended to represent a broad range of these servers/devices known in the art. Example systems suitable for use as servers **102-106** include but are not limited to servers of various capacities available from IBM of Armonk, NY, and Sun Microsystems of Menlo Park, CA. Example systems/devices suitable for use as clients **108** include but are not limited to desktop/notebook sized computers from Dell Computer of Austin Texas, and Apple Computers of Cupertino, CA, as well as various palm sized personal digital assistants from Palm Computing of San Jose, CA, and Hewlett Packard of Palo Alto, CA.

Method

[0017] Referring now to **Figure 2**, wherein a block diagram illustrating the method of the present invention, in accordance with one embodiment, is shown. As illustrated, during initialization of a directory server **102** (which is performed independent of one another), the initializing directory server **102** identifies and confirms the availability of the master and slave servers **104** and **106** available to its direction. In one embodiment, the list of nominally and non-exclusively available master and slave servers **104** and **106** are provided to an initializing directory server **102** by way of a configuration file. The configuration file includes the identities of these candidate master and slave servers **104** and **106** as well as the manner through which the master and slave servers **104** and **106** may be contacted/addressed, e.g. their IP addresses. The information may be organized in the configuration file in any convenient manner. An initializing directory server **102**

reads and processes this configuration file during initialization upon power up or reset.

[0018] In one embodiment, an initializing directory server **102** confirms the nominal non-exclusive availability of the listed master and slave servers **104-106** by “pinging” each of the listed servers with a message to solicit a response from these servers to confirm their respective presence and nominal non-exclusive availability. The “pinging” and replying may be performed in any one of a number of networking discovery protocols known in the art, including but are not limited to DHCP.

[0019] Note that the present invention provides for each initializing directory server **102** to confirm the engagement of master and slave servers **104-106** independently without consultation or coordination with other directory servers **102**, and the engagements are on a non-exclusive basis. Accordingly, some or all of the master and slave servers **104-106**, unless purposely segregated or segmented by the network administrators, may engage and “service” applications/clients under the direction of multiple directory servers **102**. However, as will be described in more detail below, the present invention provides for directory servers **102** to take into account any overlapped engagements, when performing selections and assignments of master and slave servers **104-106**. Thus, the desired ease of expansion and contraction, i.e. scalability, may be provided.

[0020] Continuing to refer to **Fig. 2**, having nominally confirmed the availability of the master and slave servers **104** and **106** available for its direction, directory servers **102** await for clients requests for accesses to the hosted application services. In due course, one or more clients **108** make the access requests, block **204**. In one

embodiment, clients **108** make the access requests by accessing web pages hosted by a front end load balancer (not shown) of directory servers **102** and selecting the appropriate options provided on the web pages. The manner in which the web pages are organized or the manner in which the options are provided thereon, as well as the manner in which the front end load balancer selects directory servers **102** are not essential to the practice of the present invention. They may be organized, provided and selected in any manner.

[0021] In response, a selected directory server **102** selects and assigns the appropriate slave server **106** to facilitate the requesting client **108** to access the request application service, and if the application service has not been instantiated, the directory server **102** also selects and assigns the appropriate master server **104** to host the requested application service, block **206**. As part of the selection and assignment process, client **108** is notified of the selected slave server **106**, in particular, the manner through which the selected slave server **106** is to be contacted, e.g. its IP address. Similarly, the selected slave server **106** is notified of the assignment of client **108**, and the master server **104** hosting the application service to be accessed by assigned client **108**, e.g. their respective IP addresses. Of course if the application service is not already hosted by master server **104**, it is also instructed to instantiate and start hosting of another instantiation or session of the requested application service.

[0022] In one embodiment, each directory server **102** performs the selections and assignments of master and slave servers **104** and **106** in accordance with a number of decision metrics that convey the ability of the non-exclusively engaged master and slave servers **104** and **106** to further undertake additional hosting of

application services or facilitating access to hosted application services, in view of the hosting and facilitating already undertaken. In one embodiment, these decision metrics include the processing resources (as measured in e.g. the amount of instructions the master/slave server **104/106** is capable of executing per time unit, the amount of memory resources available and so forth), the number of application services and/or sessions already hosting, and the number of clients already facilitating.

[0023] In one embodiment, to enable each directory server **102** to perform the selections and assignments in accordance with the decision metrics, each directory server **102** periodically polls the non-exclusively engaged master and slave servers **104** and **106** for information related to the decision metrics, block **208**. The polling and reply may be accomplished using any one of a number of messaging protocols. In one embodiment, each directory server **102** also caches the information collected. The cache data may be organized and stored in any one of a number of data structure and data storage techniques known in the art.

[0024] Back at block **206**, upon notified of their responsibilities to host application services or facilitate access to hosted application services, master and slave servers **104** and **106** host the application services and facilitate client accesses accordingly. Similarly, upon notified of slave servers **106** to contact, clients **108** access the desired application services through assigned slave servers **106** accordingly, block **210**.

[0025] Thus, it can be seen from the above description, by virtue of the substantially independent and yet indirectly coordinated direction of the master and

slave servers, under the present invention, directory, master and slave servers **102-106** may be added or removed as the number of application services to be hosted increases or decreases or as the user bases of the hosted application services increases or decreases.

Example Application – Internet Talk Radio

[0026] Referring now to **Figures 3-4**, wherein two block diagrams illustrating the system and method of **Figs. 1-2** as applied to an Internet Talk Radio application, in accordance with one embodiment, are shown. In other words, for the embodiment, at least clients **108'** are coupled to directory servers **102'** and slave servers **106'** through a public network, such as the Internet. Directory servers **102'**, master servers **104'** and slave servers **106'** may be coupled to each other through public or private networks or combination thereof. The application services hosted by master servers **104'** include hosted multi-party packet based talk radio shows (commonly referred to as Internet Talk Radio). Thus, clients **108'** may be hosts **108'a**, callers **108'b** or listeners **108'c** of the various talk radio shows of the hosted Internet Talk Radio System.

[0027] For the illustrated embodiment, as part of their responsibilities for hosting the talk radio shows of the Internet Talk Radio System, master servers **104'** are responsible for mixing the audio streams of the various talk radio shows received from callers **108'b** of the various talk radio shows through their respective slave servers **106'** for listeners **108'c** of the corresponding talk radio shows. Further, master servers **104'** are responsible for routing the mixed and non-mixed audio streams to callers **108'b** and listeners **108'c** of the various talk radio shows accordingly.

[0028] At block **402**, the initialization of directory server **102'** and the confirmation of the availability of master and slave servers **104'** and **106'** under their respective directions are performed as earlier described. Thereafter, at block **404**, client hosts **108'a** of talk shows make their respective requests to directory servers **102'** to have their respective talk shows hosted. Obviously, a talk show has to be hosted before any caller **108'b** or listener **108'c** may join the talk show. However, the different talk shows may be instantiated at different times. Some talk shows may be instantiated while others are already in progress. In like manner, client callers **108'b** and listeners **108'c** make their respective requests to directory servers **102'** to join the talk shows of interest. Both the request to host a new talk show and the request to join a talk show in progress may be made through the employment of web pages and a front end load balancer as earlier described.

[0029] At block **406**, directory servers **102'** respond to the requests to host new talk shows and requests to join talk shows of interest as earlier described. That is, directory servers **102'** select and assign slave servers **106'** to requesting clients **108'** for them to access the appropriate master servers **104'** to host/participate in the talk shows of interest. Further, if applicable, directory servers **102'** select and assign master servers **104'** to host the talk shows. Directory servers **102'** notify clients **108'**, master and slave servers **104'-106'** accordingly. Directory servers **102'** perform the selections and assignments in accordance with a number of decision metrics as earlier discovered. Further directory servers **102'** periodically poll master and slave servers **104'** and **106'** for information related to the decision metrics as earlier described, block **408**.

[0030] Upon assigned the responsibilities, at block **406**, master and slave servers **104'** and **106'** perform their respective hosting and facilitating access accordingly, whereas at block **410**, clients **108'** host, call or otherwise listen to the talk shows of interest accordingly. More specifically, client hosts **108'a** submit controls to master servers **104'** hosting their respective talk shows through their assigned slave servers **106'**. Upon receipt of these controls, slave servers **106'** route the controls to the appropriate master servers **104'**. Similarly, client callers **108'c** participate in the talk shows of interest, sending their audio streams to master servers **104'** through slave servers **106'**. Upon receipt of the audio streams of the various talk radio shows, slave servers **106'** route the audio streams of the various talk radio shows to the appropriate master servers **104'** accordingly.

[0031] For the embodiment, master servers **104'**, upon receipt of the audio streams of their hosted talk radio shows, cause the audio streams of the different callers of each hosted talk radio show to be routed to the other callers (through their respective slave servers **106'**) of the talk radio show. That is, if a talk radio show has callers A, B and C, the audio stream of caller A is routed to callers B and C, the audio stream of caller B is routed to caller C and A, and so forth. Further, master servers **104'** mix the different audio streams of a talk radio show to generate a mixed audio stream for the talk radio show, and route the mixed audio stream of the talk radio show to the client listeners **108'c** of the talk radio show.

[0032] As alluded to earlier, upon received of audio streams of the various talk radio shows from master servers **104'**, mixed or non-mixed, slave servers **106'** route the audio streams of the various talk radio shows to the client callers **108'b** and the client listeners **108'c** of the various talk radio shows as appropriate.

[0033] Thus, an Internet Radio Talk System may be advantageously hosted in scalable manner, allowing resources to be gradually increased as more shows are being hosted, or more callers/listeners are "dialing" in to the hosted shows.

Example Computer System

[0034] **Figure 5** illustrates an example computer system suitable for use as a directory server, a master server, a slave server or a client device, in accordance with one embodiment. As shown, computer system **500** includes one or more processors **502** and system memory **504**. Obviously, more processing power and memory capacity should be provided if computer system **500** is used as a directory server **102** or a master server **104**, and less processing power and memory capacity may be provided if computer system **500** is used as a slave server **106** or a client system **108**. Additionally, computer system **500** includes mass storage devices **506** (such as diskette, hard drive, CDROM and so forth), input/output devices **508** (such as keyboard, cursor control and so forth) and communication interfaces **510** (such as network interface cards, modems and so forth). The elements are coupled to each other via system bus **512**, which represents one or more buses. In the case of multiple buses, they are bridged by one or more bus bridges (not shown). Each of these elements performs its conventional functions known in the art. In particular, system memory **504** and mass storage **506** are employed to store a working copy and a permanent copy of the programming instructions implementing the directory server, master server, slave server or the client logic of the present invention. The permanent copy of the programming instructions may be loaded into mass storage **506** in the factory, or in the field, as described earlier, through a distribution medium (not shown) or through communication interface **510** (from a distribution server (not

shown). The constitution of these elements **502-512** are known, and accordingly will not be further described.

Conclusion and Epilog

[0035] Thus, an improved method and apparatus for scalably hosting application services, in particular an Internet Talk Radio System and its hosted talk radio shows, has been described. While the present invention has been described in terms of the above illustrated embodiments, those skilled in the art will recognize that the invention is not limited to the embodiments described. The present invention can be practiced with modification and alteration within the spirit and scope of the appended claims. Thus, the description is to be regarded as illustrative instead of restrictive on the present invention.